

Position Statement: Methodology to Support Dependable Survivable Cyber-Secure Infrastructures

Frederick T. Sheldon, Stephen G. Batsell
Computational Sciences and Engineering
Oak Ridge National Laboratory¹
Oak Ridge, TN 37831 USA
SheldonFT | BatsellSG@ornl.gov

Stacy J. Prowell, and Michael A. Langston
Department of Computer Science
University of Tennessee
Knoxville, TN 37996 USA
SProwell | Langston@cs.utk.edu

Abstract

*Information systems now form the backbone of nearly every government and private system. Increasingly these systems are networked together allowing for distributed operations, sharing of databases, and redundant capability. Ensuring these networks are secure, robust, and reliable is critical for the strategic and economic well being of the Nation. This paper argues in favor of a biologically inspired approach to creating survivable cyber-secure infrastructures (SCI). Our discussion employs the power transmission grid. **Keywords** Infrastructure Vulnerability, Reliability, Cyber-Security, Software Agents, Autonomic Computing Paradigm*

1 Introduction

Survivability of a system can be expressed as a combination of *reliability, availability, security, and human safety*. Each critical infrastructure (component) will stress a different combination of these four facets to ensure the proper operation of the entire system(s) in the face of threats from within (malfunctioning components, normal but complex system interrelationships that engender common failures) and threats from without (malicious attacks, and environmental insult, etc.). Structured models allow the system reliability to be derived from the reliabilities of its components. The probability that the system-of-systems survives depends upon each of the constituent components and their interrelationships as well as system-of-systems relationships. Reliability analysis provides insight to developers about inherent (and defined) components and/or (intra-)system “weaknesses” [2-5]. Naturally, as the software/system complexity increase, the reliability analysis task becomes more difficult.

To counter constantly increasing computing complexity and pervasiveness, at the core of an autonomic system is introspection and self-management. Such systems strive to (transparently) provide users with a machine/system that

runs at peak performance 24x7. Like their biological complement, autonomic systems maintain and adjust their function in the context of changing components, workloads, stress, and external conditions and in the context of hardware/software failures, random or malicious [6, 7].

1.1 Biologically Inspired Survivability Paradigm

The next generation of high performance dynamic and adaptive nonlinear networks, of which power systems are an application, will be designed and upgraded with interdisciplinary knowledge for achieving improved survivability, security, reliability, reconfigurability and efficiency². Moreover, there is an urgent need for the development of innovative methods and conceptual frameworks for analysis, planning, and operation of complex, efficient, and secure electric power networks³.

Survivable cyber-secure infrastructures (SCI) represents the combination of performance and reliability modeling, and survivability analysis germane to future (fourth-generation) power distribution and electronic information infrastructure applications including communication, network-centric distributed command and control as they relate to electrical energy generation, storage/distribution, and electrical machinery and equipment. Two important themes form the basis for increasing robustness in large-scale networked information systems. First, *cognitive immunity* promises improved, cost effective technologies for the detection, quantification and recovery from

¹ This manuscript has been authored by UT-Battelle, a contractor of the U.S. Government (USG) under Department of Energy (DOE) Contract DE-AC05-00OR22725. The USG retains a non-exclusive, royalty-free license to publish or reproduce the published form of this contribution, or allow others to do so, for U.S. Government purposes.

² A previously unknown software flaw in a widely deployed General Electric energy management system contributed to the scope of the 14 Aug. 2003 blackout (see <http://www.securityfocus.com/news/8016>). The bug in GE Energy's XA/21 system was discovered by an intensive code audit in the weeks following the blackout, according to FirstEnergy Corp., the Ohio utility where investigators say the blackout began. "It had never evidenced itself until that day," said spokesman Ralph DiNicola. "This fault was so deeply embedded, it took them weeks of pouring through millions of lines of code and data to find it."

³ The continued security of electric power networks can be compromised not only by technical breakdowns, but also by deliberate sabotage, misguided economic incentives, regulatory difficulties, the shortage of energy production and transmission facilities, as well as the lack of appropriately trained engineers, scientists and operations personnel.

vulnerabilities/faults⁴. Such cognition is truly context dependent. For example, to establish immunity in the power distribution and electronic information (PDEI) infrastructure, we must:

- Assess the state-of-practice of (remote) real-time vulnerability/fault detection for SCI.
- Use existing models of vulnerability/fault detection to develop an improved numerical simulation model for nominal/transient flows.
- Based on a new numerical model, develop and test a real-time detector that can promptly locate and accurately quantify vulnerabilities/faults.
- Simulate a real-time network of detectors and evaluate the effects of signal strength and noise on vulnerability/fault detectability.
- Explore the use of the numerical model, driven by real-time data, within a secure communications infrastructure to define the parameters of a SCI including Supervisory Control and Data Acquisition (SCADA) system.

The second theme, *self-healing*, provides biologically-inspired response strategies and proactive automatic contingency planning for the PDEI infrastructure including automated data acquisition, secure system monitoring, and control techniques between source/sink and control centers [8-10]. Realization of a self-healing prototype requires the following considerations:

- Determine the similarities between energy control (e.g., electric power grid control) and information networks for adaptation to SCI control systems,
- Assess the state-of-the-practice with respect to the application of Information Security (InfoSec) principles within existing control and information networks,
- Adapt or develop procedures for Common Mode Failure Analysis (CMFA) and Security/Survivability Systems Analysis (S/SSA) from the electric power domain to application within SCI and information networks in general (e.g., Internet) [11],
- Identify areas within SCI control and information networks where existing InfoSec technologies can be applied, but are not currently being used,
- Identify SCI specific vulnerabilities for which new InfoSec technologies and devices must be developed or adapted.

⁴ The term "fault" is used consistent with "fault-tolerant design" models, and does not necessarily refer to short circuits like "bolted faults." During the Aug. 10, 1996 west coast cascading failures one contributing cause was the McNary generator exciter circuits erroneously detecting a "phase imbalance" that was actually a drop in frequency. Frequency oscillations also contributed to voltage swings which were erroneously interpreted as "switch onto fault" logic by several protective relays that (subsequently) tripped offline. Theoretically, a fault is a discrepancy between a computed, observed, or measured value or condition and the true, specified, or theoretically correct value or condition (ANSI). Generally, a fault is an "accidental or abnormal physical" condition that *may* cause a functional unit(s) to fail to perform its required function (when and if encountered). Faults can be classified in terms of criticality indicating the severity of the failure consequences. Error analysis is the process of investigating an observed fault with the purpose of tracing the fault to its source (diagnosis).

1.2 Ensuring System Integrity

To codify biologically inspired survivability the focus should be on requirements, models and tools that aid in the process of ensuring system integrity [12] by selecting the mitigation mechanisms that maximize the individual and system wide objectives (see Fig. 1). In this way, optimization techniques can be added showing how resources can be allocated to individual solutions, and how this affects the overall survivability. An advantage of this approach, especially in the first phase, is that SCI implementations in the long haul can be targeted easier as a bottom-up approach[13]. In reality, the applicability of the such technology/ methodology to multiple energy sectors in the infrastructure scope is broad because the degree of impact (i.e., to improve or sustain energy assurance) on the energy infrastructure is determined at the component level [14, 15].

2 Network Vulnerability

Network-centric infrastructure demands robust systems that can respond automatically and dynamically to both accidental and deliberate faults. Adaptation of fault-tolerant computing techniques has made computing and information systems intrusion-tolerant and more *survivable*, but even with these advancements, a system will inevitably exhaust all resources in the face of a determined cyber adversary. Computing and information systems also have a tendency to become more fragile and susceptible to accidental faults and errors over time if manually applied maintenance or restoration routines are not regularly administered. "Fourth-generation" technologies will address these deficiencies by creating new security and survivability capabilities. Such capabilities will bring attributes of human cognition to bear on the problem of reconstituting systems that suffer the accumulated effects of imperfect software, human error, and accidental hardware faults, or the effects of a cyber attack. Vulnerabilities of particular concern include mobile/malicious code, denial-of-service attacks, and misuse and malicious insider threats, as well as accidental faults introduced by human error and the problems associated with software and hardware aging [1, 16].

The overarching goals in light of, for example, the threat posed by a blackout similar to the one that occurred on August 14 2004, is to implement systems that *always* provide critical functionality and show a *positive* trend in reliability, that exceed initial operating capability and approach a theoretical optimal performance level in the long run. Desired capabilities include self-optimization, self-diagnosis, and self-healing, and an architecture/methodology for (system-of-) systems that support self-awareness and reflection.

2.1 Survival Strategy

SCI is a strategy intended to meet the critical need for fourth generation survivability and security mechanisms that complement first-generation security mechanisms (trusted computing bases, encryption, authentication and access control), second-generation security mechanisms (boundary controllers, intrusion detection systems, public

key infrastructure, biometrics), and third-generation security and survivability mechanisms (real-time execution monitors, error detection and damage prevention, error compensation and repair). New fourth generation technologies will draw on biological metaphors (so-called artificial biology) such as software that survives because it possesses biological properties of redundancy and regeneration (i.e., parts die off without affecting the whole), natural diversity and immune systems to achieve robustness and adaptability, the structure of organisms and ecosystems to achieve scalability, and human cognitive attributes (reasoning, learning and introspection) to achieve the capacity to predict, diagnose, heal and improve the ability to provide service.

2.2 Hierarchical Evaluation

The SCI strategy uses a hierarchical method to evaluate and implement survivability mechanisms and mitigate failures associated with three important areas of energy assurance: (a) securing cyber assets, (b) modelling, and analysis to understand and enable fundamentally robust and fault-tolerant systems, and (c) systems architecture that can overcome vital limitations. Infrastructure evaluation comprises 2 phases. First, individual components of the infrastructure are evaluated in isolation to derive individual component survivability (CS). The process identifies feasible *mitigation* mechanisms on a per component basis. In the second phase, CS is composed into the system-at-large (i.e., system-of-systems). This approach leverages individual CS models to create hierarchical structures with increased system survivability (e.g., against failures due to the complexity of engaging unanticipated component interactions)⁵. To codify such an approach the focus is on models that aid in the process of ensuring system integrity [18] by selecting mitigation mechanisms that maximize individual and system wide objectives. In this way, optimization techniques can be added showing how resources may be spent on individual solutions, and consequently, how such strategies affect the overall survivability. Naturally, individual component survivability alone is not the means for understanding the survivability of the whole system-of-systems. However, using a bottom up compositional approach enables a model-based notational language to be used to provide a complete (perhaps using a hierarchical abstraction technique) and unambiguous description of the system.

2.3 Networks of Control

Industries that use and develop critical infrastructure have become more computerized, and the risk of digital disruption from a range of adversaries has increased [11]. This societal *common ground* has proven essential to our digital economy, but has become fragile and operated at its margins of efficiency without reinvestment for many years.

⁵ The sources of common mode faults are widespread. See [17] for modelling primitives that represent interdependency failures in very simple control systems (i.e., an initial step in creating a framework for analyzing reliability/survivability characteristics of infrastructures with both hardware and software controls).

Assessment and mitigation strategies are needed to support:

- Implementing and autonomously configuring optimally redundant systems,
- Low-cost data collection methodologies to create opportunities, when feasible, for *anticipatory diagnosis* of system failures,
- Identification of critically vulnerable nodes and communication pathways,
- Detecting intruders or abnormal operations, and
- Mechanisms for distributed intelligent control to effect more flexible and adaptive systems.

Fault-tolerant systems deal with accidental faults and errors while intrusion-tolerant systems cope with malicious, intentional faults caused by an intelligent adversary. Combining fault- and intrusion-tolerant technologies produces very robust and survivable systems, but these techniques depend upon resources that may eventually be depleted beyond the point required to maintain critical system functionality. A biologically inspired approach offers the ability to reconstitute and reconfigure these resources in such a manner that the systems are better protected in the process, reliability is continually improved as vulnerabilities and software bugs are discovered and fixed autonomously, and therefore the ability to provide highly available critical services is maintained.

3 Autonomic Framework

The autonomic computing (AC) approach was outlined in 2001 by Paul Horn Sr. VP of Research at IBM, as a corporate-wide initiative in response to what their customers feel are the major impediments to more widespread deployment of computing in the workplace. Customers believe that configuration management (i.e., installing software and patches, setting various performance parameters, etc) is a significant contributor to total cost of ownership.

Ideally, systems should be self-managing; work well out of the box and continue to work well as the computing environment changes (due to failure-induced outages, changes in load characteristics, addition of server capacity). New applications may be easier to deploy if existing ones can automatically adjust and if the appropriate building blocks exist to support the construction of new applications in ways that can adapt themselves. The essential theme for AC systems therefore is self-management and cognition consisting of the following four pillars [6]:

- *Self-configuration* –Automated configuration of components and systems follows high-level policies while the rest of system adjusts automatically and seamlessly.
- *Self-optimization* –Components and systems continually seek opportunities to improve their own performance and efficiency.
- *Self-healing* –System automatically detects, diagnoses, and repairs localized software and hardware problems.
- *Self-protection* –System automatically defends against malicious attacks or cascading failures and uses early warning to anticipate and prevent system wide failures.

Therefore, the higher-order cognitive processes of reflection and self-awareness are key to creating systems that are not fragile in the presence of unforeseen inputs. Moreover, these systems will have the capacity to reason, learn, and respond intelligently to things never before encountered. However, to realize the challenge many factors must be considered (see Fig. 1) and integrated into a framework (see Fig. 2).

3.1 Cognitive Cyber Defense

To achieve SCI a hierarchical method may be used to assess and implement survivability mechanisms and mitigate vulnerabilities as well as all classes of failures for the purpose of: (1) hardening cyber assets using a framework for infrastructure survivability, (2) providing robustness and fault-tolerance through modeling, simulation, and analysis, and (3) overcoming fundamental limitations for increased reliability via effective systems architecture and the application/development of the autonomic computing paradigm mentioned above [19]. Survivability assessment comprises 2 phases: First individual components of the infrastructure are evaluated in isolation to derive various components survivability. This phase identifies feasible *mitigation* mechanisms on a per component basis. In the second phase, a mapping from component survivability is extended to the overall system-of-system resulting in better comprehension that can:

- Enhance control system dependability due to fault tolerance and system integrity strategies,
- Support a modular/scaleable approach to critical systems automation, and energy/information distribution,
- Improve the ability to sustain operational capability post attack,
- Support modeling and simulation of damage phenomenology in support of more intelligent sensors, and
- Provide an optimized technology assessment approach that can be used to select system architectures and define the elements of systems and their control.

For example, using simulations run in real-time along with the controlled system to allow dynamic tuning of the simulation parameters. In this way, specific contexts can be demonstrated to show how early signs of instability can be simulated faster than real-time to predict future failures, thus offering the opportunity for *preemptive* removal of weaknesses in the control system.

3.2 Common Mode Failures

Critical energy infrastructures and essential utilities have been optimized for reliability under the assumption of a benign operating environment. Consequently, they are susceptible to cascading failures induced by relatively minor events such as weather phenomena, accidental damage to system components, and/or cyber attack. In contrast, survivable complex control structures should and could be designed to lose sizable portions of the system and still maintain essential control functions [11]. For example, in [17], the Aug. 10, 1996 cascading blackout is studied to identify and analyze common mode faults leading to the

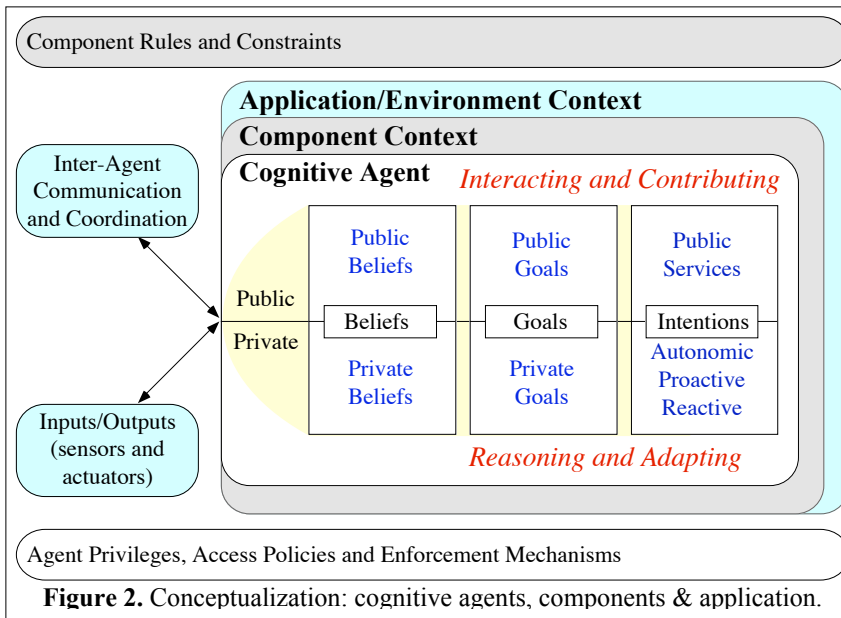


Figure 1. Automate/integrate physical, computational platform and real-world constraints to address threats and lessons learned.

cascading failure. Strategies are needed to define independent, survivable software control systems for automated regulation of critical infrastructures like electric power, telecom, and emergency communications systems.

3.3 Cyber Security

Several mitigating factors contribute to the difficulty of implementing cyber security in power substation control networks. First, is the geographic distribution of these networks, spanning hundreds of miles with network components located in isolated remote locations as well as the sheer number of devices which are connected to a single



network and open to compromise. The enormity of access points greatly increases the risk of cyber attack against electronic equipment in a substation [20].

One approach would use intelligent software agents (SAs) [21, 22](each modeled as an individual component) to deploy monitoring and control capabilities with inherent resiliency to failures [11, 23]; and desirable maintenance/evolution properties [11, 24]. SAs can enable secure, robust real-time status updates for identifying remotely accessible devices vulnerable to overload, cyber attack etc., [25, 26], as well as intelligent adaptive control (despite arguments to the contrary)[27-29].

3.4 Inherent Obstacles

The diversity of equipment and protocols used in the communication and control of power systems is staggering [11]. The diversity and lack of interoperability in these communication protocols create obstacles for anyone attempting to establish secure communication to and from a substation (or among substations in a network of heterogeneous protocols and devices). In addition to the diversity of electronic control equipment is the variety of communications media used to access this equipment. It is not uncommon to find commercial telephone lines, wireless, microwave, private fiber, and Internet connections all within substation control networks [30].

3.5 Mitigation Strategies

Previous work in this area has presented details of both threats and mitigation mechanisms for substation communication networks [30]. In [14], the most important mitigation actions that would reduce the threat of cyber intrusion are highlighted. The greatest reduction can be achieved by enacting a program of cyber security education combined with an enforced security policy. Combined, these two strategies will have the greatest impact because of the lag in cyber security knowledge within the industry.

Education and enforcement will assist with counteracting both external and insider threats ⁶.

4 Software Agents

Adaptive/intelligent software agents (SAs) [22, 23] can be used to deploy new and user-friendly data collection capabilities. SAs possess inherent resiliency to failures in responsive decision networks [31] as well as possessing software maintenance/evolution properties that promote low cost of ownership [23, 31, 32](also see [7] for a discussion of fundamental [dis-]advantages). Using software agents can enable secure and robust real-time status updates for identifying remotely accessible devices vulnerable to overload, cyber attack etc. [23, 25, 26], distributed intelligent adaptive control [29], and characterization of damage and failure mechanisms.

Cognitive systems may comprise 3 types of processes: a) *reactive*, timely response to external stimuli, b) *deliberative*, learning and reasoning, c) *reflective*, continuously monitor/adapt based on introspection.

4.1 Cognitive Agent Architecture

Based on the BDI model (see Fig. 2) [33, 34] the **Beliefs** of an agent can consist of private and public beliefs. Private beliefs represent local agent state information, which form the main basis for reasoning and reactive behavior. Public beliefs include (distributed) information about the context/environment and are the basis for reflective processes. The **Desires** are goals, where private goals govern the deliberative activities while the public goals direct the reflective processes as they describe the overall cognitive system goals. **Intentions** (services) consist of reactive, pro-active, autonomic and public plans. Reactive plans deal with timely responses to inputs and changes in the environment. Proactive plans form the core of the deliberative process, and represent the planning and reasoning processes. The autonomic plans represent the reflective process, including monitoring the agent's performance to achieve robust and secure behavior. The robustness will include, at a minimum, fail-safe plans to respond to unexpected events. An agent may make available its Public Services to other agents in the system. Agents in cognitive systems are autonomous and situated. Thus each agent is implemented with one or more active processes (or threads). For simple reactive agents (with very limited deliberative or reflective processes), a single thread is adequate to take action to various external stimuli. Complex agents may have separate threads for reactive, proactive and autonomic plans⁷.

⁶ FERC (Federal Energy Regulatory Commission) adopted NERC (North American Energy Reliability Council) security policies as standard.

⁷ Cognitive agent systems specification is defined by a Cognitive Multi-Agent Modeling Language (CMAML) and formally described using denotational semantics [6]. The key concepts of the language are Agent,

4.2 Modeling and Optimization

In addition, as an extension to the SCI, we identify how specific SCI communication protocols and mechanisms [35] can be modeled and mapped onto fault-models for understanding the impacts of common mode failures and usage profiles, including load scheduling [5, 36, 37], to identify weak points (assisting risk assessment and mitigation) in the system [17, 38, 39]. For example, there are cost effective ways to apply survivability methods [26, 40] based on redundancy and dissimilarities to the communication networks controlling the SCI. This provides several advantages: 1) the result uses a transformation model [17, 41, 42] to map the specific protocol and/or application to a graph and/or Petri Net(s) [43]; 2) interesting optimization criteria can be applied to facilitate survivability based on redundancy, while investigating the degree of independence required to achieve certain objectives (e.g., defining minimal cut sets of fault trees associated with any hazard); 3) isolation of the critical subsystems, which constitute a graph, and using agreement solutions to augment the graph to achieve the required survivability (robustness). Thus, different graphs may be derived that contain the original critical subsystems and are augmented by edges and/or vertices that allow the use of agreement algorithms. In this way, critical systems decisions are decentralized and invulnerable to malicious attacks, as long as the threshold of faulty components dictated by the agreement algorithms is not violated. Moreover, the whole field of system fault diagnosis, which originated from the PMC model (Preparata, Metz and Chien), can be applied [44, 45]. The fundamental question is "Who tests whom, and how is the test implemented to identify faulty components?" In this vein, we can specify how to derive "diagnostics" to determine if the system is robust along with a measure of confidence (i.e., see [26, 29, 40, 46, 47]).

4.3 Exemplar

Consider the need for secure web services in the context of compute-intensive applications. A natural place to focus basic research efforts is on computational problems that are hard to solve but easy to check. *NP*-complete problems are prime examples. Such a problem cannot be solved in polynomial time (assuming $P \neq NP$), and yet is easy to check, in the case of a "yes" instance due to its membership in *NP*. Within this class, let us further restrict our attention to problems that are FPT (Fixed-Parameter Tractable) [48].

A problem of size n , and parameterized by k , is FPT if it can be decided in $O(f(k)n^c)$ time, where f is an arbitrary function and c is a constant independent of both n and k . Algorithms for FPT generally operate in two stages. The first stage, termed "kernelization," is aimed at condensing an arbitrarily difficult instance into its combinatorial kernel or core. The goal is to make the kernel's size some small function of the relevant parameter (e.g., see [49]). The second stage, known as "branching," is used to explore the

search space of the kernel efficiently. It is branching that requires the vast majority of time, space and communication (e.g., see [50]). By kernelizing sequentially, but branching across the web, we achieve:

- *Verifiability*. Membership in *NP* means that we can usually expect to be able to check a candidate solution quickly. This is a critical feature, ensuring that a faulty or malicious processor cannot invalidate or subvert our computation.
- *Security*. We break the search space into disjoint sections and distribute them out to different processing elements. Each element knows only its share of the given instance, which is of course advantageous should the problem be sensitive. Even if two or more elements are untrustworthy and work in collusion, they cannot deduce the entire instance. Any attempt to exploit intercepted transmissions is similarly thwarted, thereby containing damage from intrusion. Strong concealment of the total problem is a natural part of this method.
- *Scalability*. As a computation, an FPT-based approach scales wonderfully. Branching translates to a most flexible form of partitioning. There are no *a priori* lower or upper bounds on the degree of parallelism that can be utilized. Furthermore, almost any architectural model will do, from tightly coupled parallel systems to widely distributed grids. This process can be viewed as something akin to a real-time, secure version of *seti@home* or *folding@home*.
- *Robustness*. The kernelization-plus-branching algorithm design paradigm requires no explicit communication between remote processing elements. If a limited number of elements or links fail or become unreliable, we are able to add, delete or shift branching segments around at will, thereby ensuring at worst a graceful form of degradation and preventing catastrophic failure.

5 Summary

Agent-based computing combined with the vision of autonomic computing represents an important new paradigm in Computer Science and has the potential to significantly improve the theory and the practice of modeling, designing, and implementing SCI systems. Yet, to date, there has been little systematic analysis of what makes the agent-based approach such an appealing and powerful computational model. Moreover, even less effort has been devoted to discussing the inherent disadvantages that stem from adopting an agent-oriented view. In this paper, we have sought to promote the role of agent-based software in solving complex, real-world problems of critical infrastructure protection. In particular, we argued that the development of robust, cyber defensible survivable software systems requires autonomous agents that can complete their objectives while situated in a dynamic and uncertain environment, that can engage in rich, high-level social/cooperative interactions, and that can operate within flexible organizational structures [51, 52].

Indeed, a high degree of correspondence exists between the requirements of complex system development

Belief, Goal, Plan, KQML Performative, FIPA Performative and Blackboard [34].

Cyber-Security in the Electric Sector (excerpts from [1])

The generation and delivery of electricity has been, and continues to be, a target of malicious groups and individuals intent on disrupting the electric power system. Even attacks that do not directly target the electricity sector can have disruptive effects on electricity system operations. Many malicious code attacks, by their very nature, are unbiased and tend to interfere with operations supported by vulnerable applications. One such incident occurred in January 2003, when the “Slammer” Internet worm took down monitoring computers at FirstEnergy Corporation’s idled Davis-Besse nuclear plant. A subsequent report by the North American Electric Reliability Council (NERC) concluded that, although it caused no outages, the infection blocked commands that operated other power utilities. The report, “NERC Issues Information Notice on Potential of Nuclear Power Plant Network to Worm Infection,” is available at web site <http://www.nrc.gov/reading-rm/docollections/news/2003/03-108.html> (a sobering picture of cyber-security at FirstEnergy and the affect that the “slammer” had to other SCADA systems is available at: <http://www.securityfocus.com/news/6767>).

This example, among others, highlights the increased vulnerability to disruption via cyber means faced by North America’s critical infrastructure sectors, including the energy sector. Of specific concern to the U.S. and Canadian governments are the SCADA systems, which contain computers and applications that perform a wide variety of functions across many industries. SCADA includes telemetry for status and control, as well as Energy Management Systems (EMS), protective relaying, and automatic generation control. SCADA systems were developed to maximize functionality and interoperability, with little attention given to cyber-security. These systems, many of which were intended to be isolated, are now, for a variety of business and operational reasons, either directly or indirectly connected to the Internet (e.g., in some instances, there may be a need for employees to monitor SCADA systems remotely). However, connecting SCADA systems to a remotely accessible computer network can present security risks, including compromise of sensitive operating information and unauthorized access to SCADA control mechanisms.

Security has always been a priority for the electricity sector in North America; however, it is a greater priority now than ever before. Electric system operators recognize that the threat environment is changing and that the risks are greater than in the past, and they have taken steps to improve their security postures. NERC’s Critical Infrastructure Protection Advisory Group has been examining ways to improve both the physical and cyber-security dimensions of the North American power grid. This group includes Canadian and U.S. industry experts in the areas of cyber security, physical security and operational security. The creation of a national SCADA program to improve both physical and cyber-security is under discussion in the United States and Canada.

paradigms and the key concepts and notions of agent-based computing. The agent-based computing (ABC) approach will likely succeed as a mainstream (agent oriented) software engineering ([SE] or AOSE) paradigm because it is a logical evolution from contemporary SE approaches to and because it is well suited to developing software for open systems. In contrast, ABC has the characteristic of unpredictable interactions. The strong possibility of emergent (nondeterministic) behavior in the wrong context is an inherent drawback. However, one important countermeasure is that long-term means of addressing these problems, a social level characterization of agent-based systems has been advocated as a promising point of departure. ABC should be seen in its broader context as a general-purpose model of computation that naturally encompasses autonomic distributed and concurrent systems.

6 Position Conclusions

The requirements of reliability, flexibility, and efficiency are often in conflict in large distributed control systems (e.g., SCADA systems) because the infrastructure is built and tuned independently to meet those individual requirements. Reliability requirements translate into the ability to tolerate and recover from failures and provide *a priori* (quantifiable) assurances for long-term stability. To realize a *self-healing ability*, the system must be flexible enough to dynamically adapt through reconfiguration. However, the capacity to be flexible could make the system prone to design or runtime errors and the overhead of flexibility may take away from the performance efficiency of both the control and data planes. To address these conflicting requirements at the outset, the approach must coordinate the creation and distributed layout of control

software in the form of autonomous software components or agents specifically designed to meet *a priori* service quality level needs for large complex system control⁸.

To develop more survivable distributed control architectures (i.e., SCADA), we advocate a *three-phased* approach that resolves conflicts among the different control loop performance requirements. *First*, by specifying a distributed layout of autonomous agents we can describe the end-to-end control structures at the time of system design to enable a *compile-for-service-performance* approach to the control plane. To accomplish this, a narrowly specified grammar for specifying the control framework is necessary. This language must build upon available specification methodologies such as Petri-nets and derivatives of original distributed programming/specification languages (e.g., Z, CSP, Statecharts) to enable specification of a verifiable control scheme toward gaining ultra high dependability. While formal models for distributed computation have had qualified success, the key faculty lies in translating the formalism to a network of cooperating agents. Furthermore, this step describes both the requirements and system specifications in concrete terms to enable rigorous analysis and design for (1) provisioning and resource management, (2) enabling close-to-optimal performance and (3) reconfiguration and future adaptation. Innovative graph theoretic algorithms can be used (based on formal models) to decide how to optimally

⁸ To promote and oversee the transformation of our energy system and to ensure broad, public benefit, the DOE is seeking to identify and change key regulatory barriers, catalyze open-protocol architectures and standards, develop strategic technologies when commercial interest is insufficient, conduct demonstrations to increase interest and decrease risk perceptions, and support the basic science needed to analyze and advance the transformation (see <http://www.gridwise.org>).

structure our approach: (1) reduce/abstract the size/scale of the National Power Grid problem to realistically manage the problem of validation/assessment, and, (2) make structural/architectural decisions (e.g., identify vulnerabilities/weaknesses and containment zones, as well as map agents to the grid hierarchy).

In the *second phase*, given that analytical modeling is not sufficient to accurately represent complex power grid systems, we rely on large-scale leadership class *simulation and modeling at scale* approaches to evaluate and test the deployed agent-based control scenarios. Particularly challenging at this level of complexity is the problem of faults, which originate from different sources such as hardware malfunctions and software inadequacies. In addition, faults must be minimized at the design stage and a strategy be put in place to quickly diagnose and manage dynamic faults generated during the deployments. A testing methodology that performs systematic fault coverage is lacking in the area of distributed control using agents. This is particularly true in large-scale deployments such as the power grid, and notwithstanding the existing methods, when effective, are not particularly optimized for the power grid⁹.

In the *third phase*, monitoring/control and run-time self-healing may be facilitated using *autonomous* agents having the advantage of being capable to sense/respond locally to abnormal stimuli derived from operational sensor data. An overlay network is envisioned by creating a web of communicating agents that gather and present *situational* data to higher-level control agents. Collected sensor data is stored at caching agents, and forwarded to decision centers (i.e., regional control centers) that are distributed across the transmission grid. These data sets are then correlated and fused at the centers and presented to the decision makers, either human or automated programs. By constantly monitoring the system using strategically deployed agents, problems can be quickly detected and diagnosed to activate the healing process. Self-healing networks require autonomous actuation of the network based on the dynamic sensor data to apply protective and reparative enhancements¹⁰. However, due to the complexity of the power grid the actuation of one portion of the control network can cause rolling instabilities. A contained scheme must be devised to ensure that system improvements are only made around a reliable core, whose dynamics and correlations are rigorously specified, analyzed and certified. Thus, while applied enhancements may take some time to

take effect, the reliable core ensures that the system is constantly maintained within stable operating ranges at all times.

The fundamental vulnerabilities of centralized control demand smaller, local system configurations. Resilience relies on the ability to bridge top-down and bottom-up real-time (or anticipatory) decision-making capability. The need for system integration to handle increasing complexity requires new approaches to simplify and harden the operation of complex power systems. Our energy systems require a tightly knit communications capability, whose protection will require new technology to enhance the security of command, control, and communications. Assess which are the most effective security investments because providing comprehensive physical protection to all components is not feasible. Probabilistic assessments can offer strategic guidance on where and how to deploy security resources to the greatest advantage. Ultimately, every node in the power network of the future will be awake, responsive, adaptive, price-smart, eco-sensitive, real-time, flexible, humming - and interconnected with everything else!

7 References

- [1] B. Liscouski and W. J. S. Elliott, "Causes of the August 14 Blackout in the United States and Canada," NRCAN/USDOE (US-Canada Power System Outage Task Force), Wash. DC, Interim Report, Nov., 2003, 124.
- [2] F. T. Sheldon and K. Jerath, Assessing the Effect of Failure Severity, Coincident Failures and Usage-Profiles on the Reliability of Embedded Control Systems, *ACM Symposium on Applied Computing*, Nicosia Cyprus, Mar. 14-17 2004, 826-833.
- [3] F. T. Sheldon and S. A. Greiner, Composing, Analyzing and Validating Software Models to Assess the Performability of Competing Design Candidates, *Annals of Software Engineering (Special Volume on Software Reliability, Testing and Maturity)*, vol. 8, 1999, 239-287.
- [4] F. T. Sheldon, S. Greiner, and M. Benzinger, Specification, safety and reliability analysis using Stochastic Petri Net models, *10th Int. Workshop on Software Specification and Design*, San Diego, CA, Nov. 2000, 123-132.
- [5] F. T. Sheldon, K. Jerath, and S. A. Greiner, Examining Coincident Failures and Usage-Profiles in Reliability Analysis of an Embedded Vehicle Sub-System, *Proc Ninth Int'l Conf. on Analytical and Stochastic Modeling Techniques [ASMT 2002]*, Darmstadt Germany, June 3-5 2002, 558-563.
- [6] J. O. Kephart and D. M. Chess, The Vision of Autonomic Computing, *IEEE Computer Magazine*, Jan. 2003, 41-50.
- [7] N. R. Jennings, On Agent-based Software Engineering, *Artificial Intelligence*, vol. 117 (2), 2000, 277-96.
- [8] M. Amin, Toward Self-Healing Infrastructure Systems, *IEEE Computer Magazine*, Aug. 2000, 44-53.
- [9] M. Amin, "Restructuring the Electric Enterprise: Simulating the Evolution of the Electric Power Industry with Intelligent Adaptive Agents," in *Electricity Pricing in Transition*, vol. TOPICS IN REGULATORY ECONOMICS AND POLICY Volume 42, A. Faruqui and K. Eakin, Eds. (New York: Kluwer, 2002) Chapter 3.

⁹ Simulations run in real-time along with the controlled system to allow dynamic tuning of the simulation parameters, and create opportunities, when feasible, for *anticipatory diagnosis* of system failures. Although this is a grand challenge problem in the context of high-performance simulation, specific contexts can be demonstrated to show how early signs of an instability can be simulated faster than real-time to predict future failures, thus offering the opportunity for *preemptive* removal of weaknesses in the control system.

¹⁰ The response policies and actuation techniques are driven by rule-engines at different points in the network hierarchy. These rule engines are control-system specific and communicate with the agent-infrastructure over well-defined interfaces.

- [10] M. Amin, "National Infrastructures as Complex Interactive Networks," in *Automation, Control, and Complexity: An Integrated Approach*, S. Weyrauch, Ed. (New York: John Wiley, 2000) 263-286.
- [11] F. Sheldon, T. Potok, A. Krings, and P. Oman, Critical Energy Infrastructure Survivability, Inherent Limitations, Obstacles and Mitigation Strategies, *Int'l Jr. Power and Energy Systems (Spec. Theme Blackout)*, June 2004, 86-92.
- [12] F. T. Sheldon and H. Y. Kim, Testing Software Requirements with Z and Statecharts Applied to an Embedded Control System, *Software Quality Journal*, vol. 12 (3), 2004, 232-266.
- [13] A. W. Krings, W. S. Harrison, N. Hanebutte, C. S. Taylor, M. McQueen, and S. Matthews, An Agent Supported Bottom-Up Approach to Computer and Network Survivability, *Int'l Conf. Dependable Systems and Networks (Supplement of to DSN-2001)*, Goteborg Sweden, July 1-4 2001, B70-71.
- [14] C. Taylor, P. Oman, and A. Krings, Assessing Power Substation Network Security and Survivability: A Work in Progress Report, *Proc. Int'l Conf. on Security and Management (SAM'03)*, Las Vegas, June 23-26 2003, 281-287.
- [15] H. Y. Kim, K. Jerath, and F. T. Sheldon, "Assessment of High Integrity Components for Completeness, Consistency, Fault-Tolerance and Reliability," in *Component-Based Software Quality: Methods and Techniques*, M. P. A. Cechich, and A. Vallecillo, Ed. (Heidelberg: Springer LNCS 2693, 2003) 259-86.
- [16] E. J. Lerner, "What's wrong with the Electric Grid?," *The Industrial Physicist*, Vol. 9, Issue 5, <http://www.tipmagazine.com> (Accessed: Nov. 1, 2003) Last Updated: Jan. 2004.
- [17] A. Krings and P. Oman, A Simple GSPN for Modeling Common Mode Failures in Critical Infrastructures, *HICSS-36 Minitrack on Secure and Survivable Software Systems*, Hawaii, January 2003, 334a-44.
- [18] F. T. Sheldon and H. Y. Kim, Validation of Guidance Control Software Requirements for Reliability and Fault-Tolerance, *IEEE Proc Reliability and Maintainability Symp.*, Seattle, Jan. 2002, 312-318.
- [19] C. Tristram, From Artificial Intelligence to Artificial Biology?, *Technology Review*, vol. 106 (9), Nov. 2003, 40-41.
- [20] NERC, *An Approach to Action for the Electricity Sector, Ver. 1*. Princeton, NJ: North American Electric Reliability Council) 2001.
- [21] F. T. Sheldon, M. T. Elmore, and T. E. Potok, An Ontology-Based Software Agent System Case Study, *IEEE Proc. Int'l Conf. on Information Technology: Coding & Computing*, Las Vegas, Apr. 28-30 2003, 500-06.
- [22] T. E. Potok, M. T. Elmore, J. W. Reed, and F. T. Sheldon, VIPAR: Advanced Information Agents Discovering Knowledge in an Open and Changing Environment, *Proc. 7th World Multiconf. on Systemics, Cybernetics and Informatics Spec. Session on Agent-Based Computing*, Orlando, July 27-30 2003, 28-33.
- [23] T. E. Potok, L. Phillips, R. Pollock, A. Loeb, and F. T. Sheldon, Suitability of Agent-Based Systems for Command and Control in Fault-tolerant, Safety-critical Responsive Decision Networks, *ISCA 16th Int'l Conf. on Parallel and Distributed Computer Systems (PDCS)*, Reno NV, Aug. 13-25 2003, 283-290.
- [24] F. T. Sheldon, K. Jerath, and H. Chung, Metrics for Maintainability of Class Inheritance Hierarchies, *Jr. of Software Maintenance and Evolution*, vol. 14 (3), May/June 2002, 147-160.
- [25] D. Conte de Leon, J. Alves-Foss, A. Krings, and P. Oman, Modeling Complex Control Systems to Identify Remotely Accessible Devices Vulnerable to Cyber Attack, *ACM Workshop on Scientific Aspects of Cyber Terrorism, (SACT)*, Wash. DC, November 2002.
- [26] C. Taylor, A. Krings, and J. Alves-Foss, Risk Analysis and Probabilistic Survivability Assessment (RAPSA): An Assessment Approach for Power Substation Hardening, *Proc. ACM Workshop on Scientific Aspects of Cyber Terrorism, (SACT)*, Wash. DC, November 2002.
- [27] D. P. Buse, P. Sun, Q. H. Wu, and J. Fitch, Agent-Based Substation Automation, *IEEE Power and Energy Magazine*, Mar/Apr 2003, 50-55.
- [28] J. Q. Feng, D. P. Buse, Q. H. Wu, and J. Fitch, Distributed Mobile Communication Base Station Diagnosis and Monitoring Using Multi-agents, *Intelligent Data Engineering and Automated Learning - IDEAL*, Manchester, UK, Aug. 12-14 2002, 267-272.
- [29] C. Taylor, A. Krings, W. S. Harrison, N. Hanebutte, and M. McQueen, Considering Attack Complexity: Layered Intrusion Tolerance, *Proc. DSN 2002 Workshop on Intrusion Tolerance*, June 2002.
- [30] P. Oman, E. Schweitzer, and J. Roberts, Protecting the Grid From Cyber Attack, Part II: Safeguarding IEDS, Substations and SCADA Systems, *Utility Automation*, vol. 7 (1), Jan/Feb 2002, 25-32.
- [31] F. T. Sheldon, T. E. Potok, and K. M. Kavi, Multi-Agent Systems for Knowledge Management and Decision Networks, *Informatica*, vol. 28, June 2004, 79-89.
- [32] T. E. Potok, N. D. Ivezic, and N. F. Samatova, Agent-based Architecture for Flexible Lean Cell Design, Analysis and Evaluation, *Working Conf. on Design of Info. Infrastructure Sys.*, Melbourne Australia, Nov. 15-17 2000, 181-8.
- [33] A. S. Rao and M. P. Georgeff, BDI Agents: From theory to practice, *First Int'l Conf. on Multi-Agent Systems*, San Francisco, California, June 1995, 312-319.
- [34] K. M. Kavi, M. Aborizka, and D. Kung, A framework for the design of intelligent agent based real-time systems, *Proc. 5th Int'l Conf. on Algorithms and Architectures for Parallel Processing*, Beijing, Oct. 23-25 2002, 196-200.
- [35] Z. Zhou, F. T. Sheldon, and T. E. Potok, Modeling with Stochastic Message Sequence Charts, *IIIS Proc. Int'l. Conf. on Computer, Communication and Control Technology*, Orlando, FL, July 31 - Aug. 2 2003.
- [36] A. Krings, W. Harrison, A. Azadmanesh, and M. McQueen, Scheduling Issues in Survivability Applications using Hybrid Fault Models, *To Appear Parallel Processing Letters*, 2004.
- [37] A. W. Krings, W. S. Harrison, M. H. Azadmanesh, and M. McQueen, The Impact of Hybrid Fault Models on Scheduling for Survivability, *Int'l Wkshp on Scheduling in Computer- and Manufacturing Systems, Seminar 02231, Report 343*, Schloss Dagstuhl, Germany, June 2-6 2002.
- [38] A. Krings and P. Oman, Secure and Survivable Software Systems, *IEEE Proc. HICSS-36, Minitrack on Secure and Survivable Software Systems*, Big Island, Hawaii, January 2003, 334a.
- [39] W. S. Harrison, A. Krings, N. Hanebutte, and M. McQueen, On the Performance of a Survivability Architecture for

- Networked Computing Systems, *IEEE Proc. HICSS-35*, Big Island, Hawaii, Jan. 2002, 2534 - 2542.
- [40] C. Taylor, A. Krings, W. S. Harrison, and N. Hanebutte, Merging Survivability System Analysis and Probability Risk Assessment for Survivability Analysis, *IEEE DSN 2002 Book of FastAbstracts*, June 2002.
 - [41] A. W. Krings and M. H. Azadmanesh, A Graph Based Model for Survivability Applications, *To Appear Electronic Journal of Operations Research (EJOR)*, 2004.
 - [42] A. W. Krings, Agent Survivability: An Application for Strong and Weak Chain Constrained Scheduling, *HICSS-37, Minitrack on Security and Survivability in Mobile Agent Based Distributed Systems*, Big Island, Hawaii, January 2004, To Appear.
 - [43] F. T. Sheldon, K. M. Kavi, W. W. Everett, R. Brettschneider, J. T. Yu, and R. C. Tausworthe, Reliability Measurement: From Theory to Practice, *IEEE Software*, July 1992, 13-20.
 - [44] S. Chessa and P. Santi, Comparison based system-level fault diagnosis in ad-hoc networks, *Proc. 20th IEEE Symp. on Reliable Distributed Systems*, October 2001, 257-266.
 - [45] F. P. Preparata, G. Metze, and R. T. Chien, On the Connection Assignment Problem of Diagnosable Systems, *IEEE Transactions on Computers*, vol. EC-16, Dec. 1967, 848 - 854.
 - [46] A. Krings, S. Harrison, N. Hanebutte, C. Taylor, and M. McQueen, Attack Recognition Based on Kernel Attack Signatures, *International Symposium on Information Systems and Engineering (ISE)*, Las Vegas, June 25-28 2001, 413-419.
 - [47] C. Taylor, W. Harrison, A. Krings, N. Hanebutte, and M. McQueen, Low-Level Network Attack Recognition: A Signature-Based Approach, *IEEE Proc. PDCS'2001*, Anaheim, CA, August 2001, 570-574.
 - [48] R. G. Downey and M. R. Fellows, *Parameterized Complexity*: Springer-Verlag) 1999.
 - [49] F. N. Abu-Khzam, R. L. Collins, M. R. Fellows, M. A. Langston, W. H. Suters, and C. T. Symons, Kernelization Algorithms for the Vertex Cover Problem: Theory and Experiments, *Proc. Wkshp on Algorithm Engineering and Experiments (ALENEX)*, 2004, To Appear.
 - [50] F. N. Abu-Khzam, M. A. Langston, and P. Shanbhag, Scalable Parallel Algorithms for Difficult Combinatorial Problems: A Case Study in Optimization, *Proceedings, International Conference on Parallel and Distributed Computing and Systems (PDCS)*, 2003, 563-568.
 - [51] I3P, "Cyber Security Research and Development Agenda," Dartmouth, Hanover, NH, Institute for Information Infrastructure Protection Research Agenda, 2003, 55.
 - [52] J. Cummings, "National R&D Plan for Critical, Ver. 3," Center of Technology Commercialization, US DHS Industry Forum, May 19, 2004, 28.